Confronting Power and Corporate Capture at the FAccT Conference

Meg Young
megyoung@cornell.edu
Cornell Tech
New York, NY, USA

Michael A. Katell
mkatell@turing.ac.uk
The Alan Turing Institute
London, UK

P. M. Krafft
p.krafft@arts.ac.uk
University of the Arts London
London, UK

ABSTRACT

Fields such as medicine and public health attest to deep conflict of interest concerns present when private companies fund evaluation of their own products and services. We draw on these lessons to consider corporate capture of the ACM Fairness, Accountability, and Transparency (FAccT) conference. We situate our analysis within scholarship on the entanglement of industry and academia and focus on the silences it produces in the research record. Our analysis of the institutional design at FAccT indicates the conference's neglect of those people most negatively impacted by algorithmic systems. We focus on a 2021 paper by Wilson et al., "Building and auditing fair algorithms: A case study in candidate screening" as a key example of conflicted research accepted via peer review at FAccT. We call on the conference to (1) lead on models for how to manage conflicts of interest in the field of computing beyond individual disclosure of funding sources, (2) hold space for advocates and activists able to speak directly to questions of algorithmic harm, and (3) reconstitute the conference with attention to fostering agonistic dissensus—un-making the present manufactured consensus and nurturing challenges to power. These changes will position our community to contend with the political dimensions of research on AI harms.

CCS CONCEPTS

• **Social and professional topics** → **Codes of ethics**; **Funding**; *Political speech*; *Governmental regulations.*

KEYWORDS

pymetrics, conflict of interest, corporate capture, industry engagement, research funding, agonism

ACM Reference Format:
Meg Young, Michael A. Katell, and P. M. Krafft. 2022. Confronting Power and Corporate Capture at the FAccT Conference. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22), June 21–24, 2022, Seoul, Republic of Korea.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3531146.3533194

1 INTRODUCTION

Human rights abuses, sales of surveillance systems to authoritarian regimes, environmental irresponsibility, and other scandals have fostered urgent calls for regulation of major companies like Facebook/Meta, Google/Alphabet, and Microsoft—an outcry known as the "techlash." Meanwhile, Big Tech companies are major sponsors of this preeminent conference dedicated to confronting algorithmic harms at their source—the ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT).[1] Industry co-presence and cross-pollination at FAccT is intended to bridge the gap between research findings and design choices that affect billions of people worldwide [12]. The deep contradiction of FAccT is that the technology firms responsible for some of the biggest ethical problems in the industry pay for the salaries of many of the conference organisers—not to mention the after-party hors d'oeuvres.

In this paper, we argue that the enthusiastic participation of technology firms in the ACM FAccT conference is in service of certifying their commitment to ethics [28, 50, 60, 64, 65] and add that corporate entanglement forecloses the opportunity for these firms to be confronted by narratives of the real-world consequences of AI harms (c.f. particularly Birhane et al. [11], who provide robust evidence for this neglect at both FAccT and the closely related conference on AI, Ethics, and Society; and Greene et al. [28] who find that critiques of AI/ML are co-opted by firms into a technological solutionistic, expert lens). Instead of interrogating AI from a political perspective, FAccT contributions frequently limit their gaze to the "technical" part of "sociotechnical"—the level of data, metadata, or models [24, 28]. Such abstractions deracinate harms from lived experience and sever our field from where the work is already happening: in the fight to leverage and defend against algorithmic systems on-the-ground [23]. By operating within these limits, our community misses the opportunity to coordinate the political action that the present moment demands.

We attribute this neglect—that of de-centering the realities and personal testimonies of harm—in part to conflicts of interest intrinsic to AI ethics more broadly [50, 65]. Instead of grappling with the deep contradictions within our community of practice, the conference elides them under the imprimatur of a published proceedings. This maneuver channels collective labor (including that of critical scholars intervening via Critiquing and Rethinking Accountability, Fairness, and Transparency (CRAFT) track) into an engine for manufacturing consensus behind an AI ethics that is technical, apolitical, documentation-oriented, and unregulated [24, 28]. In this regard, AI ethics joins the ranks of canonical examples of conflict of interest such as public health research funded by big tobacco, pain research funded by opioid manufacturers, or climate research funded by oil companies [1]. In each case, industry actors instrumented research inquiry, methods, findings, and policy implications toward those favorable to their bottom lines.

[1] "ACM FAccT Sponsors and Supporters," ACM FAccT Conference https://facctconference.org/2021/sponsorship.html.

FAccT '22, June 21–24, 2022, Seoul, Republic of Korea

ACM ISBN 978-1-4503-9352-2/22/06.
https://doi.org/10.1145/3531146.3533194

This paper brings the unease expressed between colleagues to bear on corporate entanglement at FAccT and underlines the consequences of cleaving ethics from justice. We call on FAccT to acknowledge the political dimension of its work toward its stated aim of advancing more ethical AI. In particular, we argue that (1) FAccT should model itself on the medical research field, which promotes fuller data collection and disclosure of researcher interests, (2) center the critical contributions of advocates and activists from outside predominantly white institutions in the Global North, and (3) turn towards a convening that fosters agonistic dissensus to challenge—rather than certify—those in power.

2 CORPORATE CAPTURE OF THE ACADEMY AND CONFLICTED RESEARCH

A longstanding body of literature describes the enmeshment of industry and academic research. A particular focus is on the modern university, which is described as an expression of *neoliberalism*: an ideology that embraces market-based solutions and a diminished role for both government intervention and community oversight. Whereas scholars like Shapin have argued that academic research has always been inflected by private interests [57], Mirowski marks a shift from the values and policy imperatives that supported scientific practice through the Cold War period to a wider cultural embrace of a neoliberal ideology that dominated university and corporate research in turn [47]. This era transformed the perceived value of academic research from a public good into a private commodity, resulting in the encroachment of a privatization regime further into the university in the form of patents over faculty discoveries, restrictive contracts around proprietary lab equipment, and the growing role of startups and contractors in scientific discovery [47]. Where Mirowski attributes this shift to the influence of neoliberal economists and the state's retreat from its management role in scientific exploration, Vinsel finds its roots in the opportunistic pursuit of new sources of funding by academic and industry scientists, university administrators, and firms [62].

2.1 The neoliberal academy reproduces power relations

A central vein of scholarship on the neoliberal university examines its relationship to minoritized people. Author la paperson [40] traces the settler colonial composition of American universities from their origins as land-grant institutions and the tactics they draw from colonial schooling; la paperson maintains focus on how such institutions are thwarted and re-purposed from within. Ferguson [16] describes the role the university has played in generating categories of minority difference; he finds that the 1960s student race and social justice movements that challenged academia were, in turn, subsumed by it. Drawing on a Foucauldian lens, Ferguson argues that academia is a site both of direct discrimination and of the reproduction of racial categories conducive to its goals. The co-opting of struggles for social justice produces performative diversity, equity, and inclusion initiatives [3] that extend, rather than meaningfully alter, state and corporate power by restricting the scope of what is imagined to be possible.

Foundational work by Harney and Moten argues that the university is inextricable from state and corporate power [33]; they focus particularly on the university's role in professionalization [33, 46]. Those who deviate from this social reproduction are deemed "unprofessional." Harney and Moten argue that professionalization operates definitionally as a form of negligence, by instructing professionals to disregard anyone who is (by another name) Other—described as "the outcast mass intellectuality of the undercommons" [33, p. 33]. Harney and Moten use the term "undercommons" to refer to a generative, outlawed, unassimilated underground of people in the university—one which is already present outside of predominantly white institutions. In response to the question of challenging the corporate capture of academia, Harney and Moten find that scholars who purport to be critical of the neoliberal university do not pose a threat to it. Rather, academic criticism recognizes and legitimates the university, reconstituting and completing the state's strategy, thereby continuing the survival of the institution in its present form. Critical scholars, too, participate in neglect of the undercommons.

2.2 Conflicts of interest in academic research

A different tradition of research frames the question of industry entanglement in research as *conflict of interest*. A number of academic fields provide instructive case studies in this regard. Oreskes and Conway document how the tobacco industry shaped 20th century health science in its favor [51]. The authors use the term "the tobacco strategy" to describe how decades of diffuse, non-directive funding was able to obfuscate evidence linking smoking to health problems. The tobacco strategy was to direct funds towards scientists whose agendas were non-threatening to industry objectives, such as research examining the link between heredity and cancer. Industry lawyers cited these findings in turn as independent evidence of uncertainty in the relationship between smoking and the onset of lung disease. Oreskes and Conway emphasize how 60 years of obfuscatory science concealed harms from the public and—most critically—delayed regulatory action that would have saved lives.

Previous work in the biomedical field also evinces a strong relationship between industry funding and the substance of research findings. Sismondo [59] examines a broad array of previously published review research about the impact of pharmaceutical company funding on clinical trials, finding conclusive evidence that company sponsorship plays a significant role in the publication of results favorable to sponsors' interests. Sismondo points out that sponsored research can be methodologically rigorous, but selective publication, biased interpretation [5, 56], and other choices [7] together produce a corpus of results favoring companies' interests. The case of industry influence in biomedical research is notable for the substantial body of evidence available, which includes systematic surveys spanning quantitative analyses published between 1966-2002 [43] and between 1980-2002 [7]. The same association between sponsorship and favorable research results is evident in a wide number of domain-specific review studies (e.g. psychiatry [38], pulmonology [44], and cardiology [55]). The vastly different study conditions available in the medical field allows for detailed analysis—such as how favorable results rise in relation to the *proportion* of industry funding [5, 55]. Notably for our purposes, two studies of the contents of annual professional society meetings find that *all* industry-sponsored research contributions were favorable

to the sponsor (c.f. [19, 22]). At FAccT, similar practices by tech companies were laid bare through the events surrounding the publication of Gebru et al.'s "On the Dangers of Stochastic Parrots" [8], wherein Google aimed to censor research critical of their own business.[2]

2.3 Conflicts of interest in Big Tech research

Previous work has raised similar concerns about tech industry funding relationships [1, 50, 65]. A large majority of AI ethics research faculty at top-ranked American computer science departments have received funding from Big Tech [1].[3] At stake is how academic work is channeled into policy and practice. Tech giants have previously cited funded researchers as independent authorities—such as when then-Twitter CEO Jack Dorsey provided testimony to the U.S. House of Representatives referring to research from Cortico and an MIT lab as independent evidence, when Twitter had been closely collaborating with them both [29]; Facebook has been documented engaging in similar practices [1].

These concerns about the soundness of corporate-sponsored academic research are reprised in growing critique of Big Tech firms' societal impact.[4] In the face of rising public scrutiny, critics have described tech ethics practices as a form of white-washing [9, 26, 45, 60, 61] that serves to preempt criticism and diminish the likelihood of government regulation [64]. Rather than check corporate power, AI ethics functions to "co-opt the language and performative function of ethics to pursue self-promotional goals" [9] and to reinforce white supremacy [2, 31]. Previous attempts to critique, re-direct, or refuse the way Big Tech dominates computing scholarship are often co-opted; as Greene, Hoffmann, and Stark ask:

> "What if, instead of being brushed aside, our critiques are being heard but transformed into something we might not recognize?" [28].

Indeed, Big Tech diagnoses problems as purely technical, rather than a source for policy action [28]; Ganesh and Moss find that harms identified by marginalized groups are predominately subsumed into design fixes [23]. Where tech companies are enthusiastic about regulation, they set out to exert control over its shape [50]; companies have expended huge resources to this end [66].

2.4 Organizing in response to Big Tech capture

A number of scholars have built on advocate and activist pressure to challenge industry presence in academic conferences. In 2018, a coalition of 15 immigrants' rights groups, tech-worker and labor groups, and community organizations protested the $39 million contract between surveillance firm Palantir and the U.S. Immigration Customs Enforcement (ICE) office to maintain a database of immigrants and their relationships.[5] Drawing on the momentum from this action while also acknowledging years of prior expressions of concern, academic researchers from the Data Justice Lab at the University of Cardiff and members of DATACTIVE from the University of Amsterdam wrote a public letter withdrawing from that year's Amsterdam Privacy Conference sponsored by Palantir.[6] This action further catalyzed the community to begin a campaign called Funding Matters demanding that the Amsterdam Privacy Conference decline corporate sponsorship from Palantir and calling on computing conferences to "develop sponsorship criteria and guidelines that ensure academic independence and proper consideration of human rights."[7] The Funding Matters campaign launched simultaneously to the immigrant rights group Mijente's now long-standing NoTechForICE campaign.[8] The following year, U.S.-based colleagues called on the Privacy Law Scholars Conference to drop Palantir as a sponsor, citing its role in "building technologies that support federal immigration enforcement policies to profile and deport immigrants, detain children, prosecute families, and conduct surveillance on low-income communities."[9] A panel and campaign at the Computer-Supported Collaborative Work Conference was inspired by these events to discuss academic integrity in industry funding of HCI [29].[10]

These concerns have also been explicitly directed toward FAccT as a conference. In 2021, a group of concerned members of the FAccT community came together to express concern about the lack of mandatory funding disclosures by conference authors, the inclusion of work involving controversial technologies without meaningful political analyses, the production of algorithmic audits sponsored and co-authored by employees at audited firms, and the over-representation in the conference of researchers from predominantly white institutions in the Global North [36]. Published in 2022, an anonymous survey of 60 FAccT authors, reviewers, and organizers also gave voice to concerns about corporate conflict of interest: Laufer et al. [41] report:

> One participant warned that "the FAccT community should be more careful in how researchers from big tech industry are assigned key positions in the conference." Another described the potential harms of unregulated corporate influence as follows: "It is easy for industry to adopt (or co-opt) some aspects of the work in ways which only minimally help those affected by algorithmic systems, and may even make them worse by giving them a false sense of legitimacy." Corporate interests can additionally bias the scholarly discourse: "Some questions are not asked or are

[2] T. Simonite, "What Really Happened When Google Ousted Timnit Gebru," *Wired,* Jun. 8, 2021. https://www.wired.com/story/google-timnit-gebru-ai-what-really-happened/.

[3] Of course this problem is not unique to AI ethics, for instance with the MIT College of Computing never having responded to the Movement for Anti-Oppressive Computing Practices demand to account for the College Dean's conflict of interest sitting on ICE contractor Amazon.com's Board of Directors. "Update: MIT College of Computing Still Must Address Systemic Racism" Movement for Anti-Oppressive Computing Practices, *Medium,* Jul. 12, 2020. https://medium.com/@macprac/update-mit-college-of-computing-still-must-address-systemic-racism-51c9f9f3ec9d.

[4] R. Neslage, "Techlash? America's Growing Concern With Major Technology Companies," *Knight Foundation,* Mar. 11, 2020. https://knightfoundation.org/reports/techlash-americas-growing-concern-with-major-technology-companies/.

[5] S. Dremman, "News Protesters demand Palantir end ICE contract," Palo Alto Weekly, Jul. 31, 2018. https://www.paloaltoonline.com/news/2018/07/31/protesters-demand-palantir-end-ice-contracts.

[6] S. Milan, "Why we won't be at APC 2018," DATACTIVE Project, University of Amsterdam, Sept. 19, 2018. https://data-activism.net/2018/09/why-we-wont-be-at-apc-2018/.

[7] "APC2018 / Palantir," Funding Matters. https://fundingmatters.tech/statement/.

[8] No Tech for ICE campaign https://notechforice.com/.

[9] "Scholars tell UC Berkeley: Cut ties with Palantir," *Mijente,* May 29, 2019. https://mijente.net/2019/05/scholars-tell-uc-berkeley-cut-ties-with-palantir/.

[10] Both the CSCW panel, "Patron or Poison?: Industry Funding of HCI Research" and campaign were organized by our team.

> difficult to ask within industrial research. There is a political bias."

Laufer et al. describe concerns such as these as a major theme emergent from their questionnaire responses, attesting to a shared focus within our community on how to manage these relationships.

3 HOW FACCT MANUFACTURES CONSENSUS

Having described pitfalls of financial incentives from technology firms and sponsored academic researchers, we present a case study of a paper accepted and published by FAccT in 2021 as an example of conflicted work within its proceedings. We find that acceptance into the conference ratifies the contents of the paper to society at large and draw out the consequences of that ratification. We provide other examples where research integrity at FAccT can be perceived as impacted by corporate entanglement, then explore how these problems are related to the conference's design.

3.1 Publishing and certifying corporate apologia

A 2021 FAccT paper, "Building and Auditing Fair Algorithms: A Case Study in Candidate Screening", describes an analysis of an algorithmic employment screening and selection system produced by the company, pymetrics (sic, hereafter capitalized) [67]. Wilson and colleagues conducted a "cooperative audit," described in their paper as a "framework for external algorithm auditors to audit the systems of willing private companies" (p. 666). According to Wilson et al., Pymetrics granted the researchers access to aspects of its system's functionality that are otherwise made inaccessible to external inquiry and which are closely guarded as intellectual property. In addition to granting permission and providing access to source code and model design details, Wilson's research team at Northeastern University received funding from Pymetrics and agreed to limit the scope of their analysis to a set of specific criteria given by the company. Four of the paper's eight authors are Pymetrics employees, including its CEO as last author—a significant position in an author list which is often (though not always) reserved in computer science papers for senior or supervising authors.

3.1.1 The Pymetrics product. Pymetrics provides a job applicant screening product that performs behavioral assessments using an ML-based predictive model and evaluates how applicants perform on proprietary tasks similar to IQ test questions. The model is trained by Pymetrics engineers using data collected about test answers of employers' nominally "successful" incumbent workers. The trained model analyzes the responses of job applicants and produces scores that predict their similarity to incumbent workers.

Like other companies that produce algorithmic evaluations of prospective employees, Pymetrics channels employers' desire to increase efficiency in hiring. A major selling point is that their systems reduce "human bias," particularly by meeting employers' desire to reduce their exposure to actionable charges of employment discrimination. Marketing materials produced by Pymetrics and those produced by other companies offering algorithmic employee selection tools, such as HireVue, suggest that their products are tools for addressing employment discrimination.[11] Following on this market positioning, Pymetrics calibrates their model to meet a specific employment discrimination standard. In the United States, legal protections for job applicants against adverse impact in employee selection are found in the Uniform Guidelines on Employee Selection Procedures (1978), which defines "adverse impact" as a "selection rate for any race, sex, or ethnic group which is less than four-fifths (4/5) (or eighty percent) of the rate for the group with the highest rate..." The highest rate is generally assumed to be cisgender men of European descent but includes anyone who does not fall into a category described by one of the legally protected classes. Under US law, legally protected classes are age, disability, gender reassignment, race, religion or belief, sex, sexual orientation, marriage and civil partnership, and pregnancy and maternity. Pymetrics claims that internal validation of their model based on the four-fifths standard indicates that the system is "fair". Wilson and colleagues report in their paper that the tool met this standard during their analysis by exhibiting no more than a 20% bias against applicants from any single protected demographic category.

3.1.2 Conflict of Interest Concerns. We find several conflict of interest concerns in this published research. First, there is a strong appearance of a conflict between the research team and the firm. The inclusion of Pymetrics employees as authors of "Building and Auditing Fair Algorithms: A Case Study in Candidate Screening" raises questions about the editorial independence of the work. Beyond the authorship is the financial relationship. The company provided funding to Wilson and his colleagues at Northeastern University and the scope of the audit was defined by the company, precluding various scenarios from scrutiny. The authors were open about this relationship, their funding, and the constraints they agreed to as disclosed in the paper—although it is unclear to what extent this disclosure was evident to the paper's reviewers during double-blind peer review. Regardless, the authors' candidness in the final published version of their paper does not resolve our questions about the independence of the research. While the paper's positive conclusions, which support the claims made by Pymetrics, may have been the end result under a different set of arrangements, it seems unlikely that this paper would have been submitted had the results negated the firm's validation claims. Even the appearance of a conflict of interest should raise questions for conference organizers and members of the research community as a matter of reputation and credibility. In addition to the disclosed direction-setting by Pymetrics in laying out the terms of the research, the question of independence is challenged by the decision to include Pymetrics employees—including its CEO, a prominent lobbyist for the company's interests—as co-authors on a paper concerning their product.

3.1.3 Black Boxing Employment Discrimination. Second, algorithmic job selection is a controversial topic. The concept of "fairness" in this arena is contested; due care must be used in its definition as its use in the employment context can have particularly high stakes for prospective job applicants. As Pymetrics accurately states

[11] C.f. "Removing Bias from Talent Decisions with Artificial Intelligence," *Pymetrics*, n.d. http://go2.pymetrics.ai/removing-bias-from-talent-decisions-asset-download-ty and "Find the right match between candidate potential and need," *Hirevue.* andhttps://www.hirevue.com/why-hirevue.

on its website, "a well paying job is a lifeline to economic opportunity".[12] Labor relations are a longstanding site of struggle for economic and social justice; discrimination in hiring is major target. Decades of activism and legal action regarding employee selection have produced legal frameworks in many countries intended to blunt the worst historical abuses in hiring practices, but this has not eliminated employment discrimination. In US law, it is up to workers to prove a discrimination claim, which is a high burden many cannot overcome in practice due to the requirements for proving a case [39]. Many workers continue to struggle to have legitimate claims recognized. This is especially so for workers whose identities and abilities do not fit cleanly into the classes articulated in anti-discrimination law.

Employers have vested interests in the worker-employer relation but enjoy some advantages in defending themselves—particularly large employers—including access to greater financial resources, retained legal teams, and increasingly advanced technologies for sorting and managing workers. There is little evidence that the introduction of data-driven technologies into this relation improves conditions or levels the playing field for workers; instead, it is more likely to provide benefits predominantly to employers. Products like Pymetrics are useful in part as a liability shield against discrimination claims; Wilson et al. missed the opportunity to more fully engage with such questions of power, law, and politics in the role of technology in the fraught relations of employment.

3.1.4 Manufacturing Consensus in Defining Fairness. From what is evident in the published paper, Wilson et al. did not challenge the definition of fairness provided by Pymetrics. The study instead aims to measure the Pymetrics system's selection outputs tabulated against the UGESP standard of adverse impact. This is a very narrow account. Scholarship on algorithmic fairness has shown that fairness is a rich and complex concept subject to dimensions of arbitrary narrowing that both obscure and reproduce structures of social injustice [34]. A more thorough account would build on traditions from human computer interaction, like value sensitive design, which acknowledge competing interests in these design decisions [21] and call for greater accountability from those who produce and adopt technologies that intervene in human lives [13]. Philosophers of technology and legal scholars have also appealed for standards of accountability and that access to recourse be provided to persons affected by algorithmic decisions based on unfavorable results [10, 63]. Where scholarship on algorithmic fairness often favors a powerful actor [60], the paper by Wilson and colleagues is an example of treating fairness with a casual violence that leaves legitimate questions unaddressed.

3.1.5 Conflicts of Interest Neglected by FAccT Peer Review. Many of the conflicted aspects of the paper by Wilson et al. were likely unknown to reviewers due to the anonymization of its authorship.[13] Even so, several clues were likely present in its double-blind version. In the published version of the paper, the authors are clear about the types of analyses they did not perform and the criteria they were given; the work is also declared to be a cooperative venture between the firm and the researchers, which included a contractual relationship between the researchers' academic institution and the firm. The paper indicates that "Pymetrics contributed to the background, design, and discussion of this manuscript, but did not materially author nor alter any audit methods or results of the audit." (p. 667) While some of these facts may have been included in the reviewed version, they were unlikely to be readily apparent without the list of co-authors that includes company employees and the company CEO. Language in the final paper minimizes these relationships, as the paper states that the audit was conducted "at arms length" from the company (p. 675) and also that they were granted complete independence. Full funding disclosures were not mandated by FAccT during the review stage in 2021, such that reviewers may not have been aware of the funding provided by Pymetrics.

3.1.6 Pymetrics Marketing its Product as "Independently Audited". The potential downstream effects of the FAccT conference's inattention to this issue is illustrated by the paper's acceptance at FAccT having been later leveraged by Pymetrics in their marketing and lobbying efforts. The company now claims that their product has been "independently audited" in its marketing materials.[14] We note that the terms "audit" and "external" lack definitional precision; scholars seeking to develop a definition of an algorithm audit include Raji et al. [54], who define an internal audit as one performed by company employees or contractors while an external audit is one performed without privileged access to the system, and by experts without formal association with the producer of the system. Wilson et al. (including the Pymetrics co-authors) cite these definitions by Raji et al. in their paper, which undermine classification of the audit as external. Nevertheless, Pymetrics uses that term to describe the analysis by Wilson et al.

3.2 Other Reflections on FAccT as a Conflicted Body

Several other facets of the FAccT conference demonstrate its entanglement with industry: its luxurious conference events, program committee memberships, conference sponsorships, and closed access to research data.

3.2.1 Luxurious industry-sponsored conference parties. Attendees at in-person FAccT conferences are typically offered access to networking receptions underwritten by technology companies. In the 2020 FAccT conference city of Barcelona, Spain, Facebook hosted a well-attended reception replete with fine food and catering in a warmly-lit restaurant resembling a wine cave. That same year, Google's AI Ethics team rented a cavernous, light-drenched, exposed-brick flower shop outside the conference hall to host a notable CRAFT workshop called the CtrlZ.AI Zine Fair [32].[15] While

[12] "Soft skills platform," *Pymetrics,* n.d. https://www.pymetrics.ai/mission

[13] At the time of writing, the ACM FAccT conference has been using a double-blind peer review process to evaluate submitted manuscripts for inclusion in the conference proceedings. "ACM FAccT 2021 Call for Papers", https://facctconference.org/2021/cfp.html

[14] K. Trindel, "Third-Party Auditing: Why We Invited External Experts into Our Codebase" *Pymetrics,* May 13, 2020. https://www.pymetrics.ai/pygest/third-party-auditing-why-we-invited-external-experts-into-our-codebase.

[15] The CtrlZ.AI zine fair is the most successful example of FAccT featuring activist perspectives to date—even as it occurred concurrently outside the conference venue. The day-long event hosted tables and workshops from a number of groups including the Stop LAPD Spying Coalition, CodingRights, Free Radicals, the Coveillance Collective, and the Design Justice Network. In addition to its work to center activist voices, this

industry-sponsored events are commonplace at academic conferences, their role at FAccT reinforces the suggestion that industry presence is core to this community.

3.2.2 Industry representation among conveners. Industry-affiliated scholars occupy important organizing roles and other roles of responsibility at the FAccT conference. At least one member of the current Executive Committee and at least nine members of the Steering Committee list their affiliations as technology firms, including Google, IBM, Microsoft, and Sony. As is commonplace in our field, other committee members have significant ties to industry in the form of major grants or past employment in the sector. We note, for example, that Wilson was a member of the FAccT Executive Committee when the Pymetrics paper was accepted to the conference. Conference organizers past and present have also worked directly with law enforcement agencies and the military even while primarily affiliated with a university or other academic institution. These relationships position organizers to set the direction of FAccT even while many technologies subject to the analysis of its authors are potentially damaging to the interests of communities who are not as well represented. Notably, there are few organizers listed as primarily affiliated with civil rights or activist organizations; there is no one from such groups currently listed as a member of the Executive Committee.

3.2.3 FAccT conference sponsorship. The primary sponsors and supporters of the FAccT conference are companies. This sponsorship is subject to rules laid out by the conference's sponsorship policy. In 2021, FAccT was sponsored by Google DeepMind, Facebook AI, the Alfred P. Sloan Foundation, Twitter, IBM, Microsoft, the Vector Institute, the MacArthur Foundation, the Ford Foundation, OpenAI, and Luminate.[16] Previous years had a similar makeup: in 2019 the conference was sponsored by DeepMind, Google, Microsoft, Spotify, Twitter, Knight Foundation, Ford Foundation, Luminate, MacArthur Foundation, Open Society Foundations, and Schmidt Futures.[17]

3.2.4 Proprietary research data access. Non-disclosure agreements present a barrier to what information is available to a researcher and what becomes possible to publish. Firms that subject employees to NDAs make it more difficult for external researchers have visibility into their operations, to reach and recruit participants, garner interview agreement and approval, or to publish results. A commonly used technique for HCI researchers to gain access to a firm is to themselves sign a non-disclosure agreement in order to get "behind the curtain," while limiting what might be reported from their findings. NDAs create and replicate silences in the research record of industry software development practices.

intervention helped to lend Google a radical valence among conference attendees. For photos, see CtrlZ.AI Zine Fair https://twitter.com/ctrlz_ai_zines.

[16] ACM FAccT Sponsors and Supporters https://facctconference.org/2021/sponsorship.html. See also the ACM FAccT 2021 Sponsorship Policy https://facctconference.org/2021/sponsorship_policy.html.

[17] "ACM FAT Sponsors and Supporters," ACM FAccT, 2019. https://facctconference.org/2019/sponsors.html.

3.3 FAccT forecloses critique

The current institutional design of FAccT glosses industry-funded studies alongside and under the same banner as more critical research. This practice manufactures the appearance of consensus by our broader community in presenting all published proceedings from differently situated interests as of like kind and authority. It allows FAccT to present as something other than an industry-supported conference. Meanwhile, the scope of inquiry is largely characterized by a lack of engagement with the political nature of knowledge production, particularly in a field as awash with funding and market opportunities as AI (and computing more broadly). At stake is the FAccT community legitimizing corporate perspectives about who is responsible for AI harms, rather than independently assessing them. Indeed, critical scholars who commit resources to FAccT through publishing, reviewing, and organising are obliged through this institutional design to allow their participation to endorse—even unknowingly—research they would not otherwise. Participating in FAccT therefore further entrenches corporate power in computing and in broader settings such as law and policy.

4 RECASTING FACCT IN A DIFFERENT MOULD

Corporate capture of the conference has foreclosed the possibility of FAccT being a space that is welcoming to critical voices. Their absence positions corporations and academics paternalistically as problem-solvers for—rather co-organizers with—those harmed by software systems. We respond through a call to change the institutional design of the FAccT conference so that vigorous critique, wide-ranging viewpoints, and testimonies of personal experiences with harms of technology are not limited to the conference's margins. First, the conference can draw on practices from medical research to mandate funding disclosures—a change that we note to be already in progress. We call on FAccT to go beyond disclosure and lead the ACM in developing standards and tools for wider community awareness of the funding landscape, such as a field-wide registry of research funding. Second, the conference can become a space for a substantially greater number of people, even a majority of people, drawn from advocacy and activist roles, while aligning the benefits of conference participation more closely with those participants' theories of change. We invite the FAccT organizing committee to engage in a formal outreach process to a wide range of activists, advocates, and community organizers—such as through paid focus groups—to elicit what benefits of conference participation could look like from their perspectives. Third, we identify ways that FAccT can draw on the political philosophy of agonism [49] to scaffold confrontations and creative dissensus, rather than gloss over these differences under the banner of a single conference proceedings.

4.1 Locating our funding and ourselves

As long as FAccT is funded and led by tech companies, it is not well-positioned to critically examine tech products and practices. In disciplines like medicine, we note a wide recognition of concerns about conflict of interest (COI). A committee convened by the National Academies on this topic published a set of helpful guidelines

for identifying, assessing, and managing conflicts of interest in clinical medicine [18]. They note that even when conflicts are disclosed to the public, they can still remain unmanaged; managing conflicts of interest can include disclosure, independent review, insulating students, and divesting from equity. The report advises that a sufficient amount of detail should be provided for the public to assess the risks of COI. Importantly, they acknowledge that COI concerns are not limited to research integrity alone—recall the aforementioned systematic literature reviews which found that industry-funded research can meet a high standard methodologically but become skewed in data interpretation and framing [5, 56, 59]. Beyond research integrity concerns, they highlight the risk unmanaged conflicts pose to public trust.

Guidelines from medicine further highlight the complexity of funding disclosures and the high burden of time and energy involved. Many of us working in academic research are obliged to negotiate funding relationships as a matter of institutional survival. Funding sources are often pooled at the level of a lab, department, or institution such that even individual academic workers may have limited visibility into the sources sustaining their own work (particularly student workers).[18] To allay the burden of this work, the National Academies call for their field to provide greater standardization and infrastructure for tracking these relationships:

> Such standardization is best pursued through a consensus development process that involves a broad array of concerned parties (e.g., academic medical centers, professional societies, public interest groups, and NIH and other public agencies). On the basis of the agreements resulting from this process, the next step would be for software developers to produce computer programs that allow an individual to fill out a standard questionnaire and then format the information for different institutions and purposes.

A standardized process like that suggested by the National Academies would be an important addition to fields like computer and information science, where ties to technology firms are widely considered to be important channels for data and platform access, professional development, and real-world impact.

We call on FAccT to lead the way within the ACM for such a standardization and data collection process. [19] A standardized process and an ACM funding explorer will help illuminate the state of the funding landscape in our field; it can also reduce the risks to individual researchers of full transparency through collective action—where we currently rely on individual responsibility in an honor system for self-assessment. We imagine a longitudinal reporting process interoperable between institutions akin to the unique Open Researcher and Contributor IDentifier (ORCID) tool; this process would collect longitudinal data on funding at the level of labs, departments, and institutions toward a fine-grained data source for rendering our funding and affiliations more legible. The confidentiality and publishing of this data set could be managed by data stewards within the ACM (c.f. [69]); interactive data visualizations of its findings would be valuable additions to the FAccT or ACM websites. It would also support the community to collectively reflect on our funding landscape, which previous work notes is heavily skewed toward industry sources given the lack of readily available public funding at this historical moment [1, 29].

4.2 Sharing space with advocates, activists, and interested publics

Throughout the world, governments and other powerful actors enact authoritarian agendas through their use of artificial intelligence, machine learning, and a range of communications platforms that rely upon these technologies. For example, the Bharatiya Janata Party currently leading the government in India employs social media platforms to foment political repression, politically motivated misinformation, and mob violence. Meanwhile, activists attempting to promote social and environmental justice using social media within such regimes face violence and repression. Climate activist Disha Ravi was arrested by the Indian government on February 13, 2021 for sedition after she shared organizing materials on social media as part of a global climate action campaign and its common cause with national farmers' protests. Ravi's arrest is one in a growing tide of others amid state-sanctioned political repression against dissidents and Muslims. While presently beyond the scope of most AI ethics scholarship, these harms were subtended and intensified by the ruling party's use of systems[20] like Twitter [4, 6, 52], Whatsapp [25], and Facebook.[21][22] AI-powered misinformation and harassment is paradigmatic of the political consequences of AI-powered systems for individuals, social movements, and societal violence.

We argue that scholarship and other contributions reflecting direct experience of violence and harm should be central to a conference dedicated to understanding and addressing algorithmic harms; particularly as some of these harms become evident only at scale [58]. FAccT to date has addressed itself primarily to people from elite academic institutions from the small footprint of countries where technology firms are headquartered.[23] Meanwhile, the highest stakes scenarios in AI ethics often lie outside these countries

[18] We also note the complexity of COI beyond the question of financial conflict; a small field like ours invites other forms of industry pressure such as navigating social ties to companies or concern about future employability.

[19] We note this year's addition of finer-grained funding disclosures to be attached to individual contributions, available at "Disclosure of Funding Sources at ACM FAccT 2022," ACM FAccT, https://facctconference.org/2022/funding_sources_disclosure.html. An optional, deeper disclosure prompt lists every stage of the research and publication process and asks about the role an employer, funder, or sponsor played in each step. Considering the burden of time and effort on individual researchers, we call for further action that would acknowledge the institutional level at which decisions happen about funding and a funding disclosure mechanism that leverages this reality towards a comprehensive approach.

[20] "India's Hindu nationalist BJP leads in disinformation race," Agence France Presse for The Express Tribune, Mar. 11, 2022. https://tribune.com.pk/story/2347475/indias-hindu-nationalist-bjp-leads-in-disinformation-race.

[21] K. Sambhav and N. Ranganathan, "Facebook charged BJP less for India election ads than others," Al Jazeera, Mar. 16, 2022. https://www.aljazeera.com/economy/2022/3/16/facebook-charged-bjp-lower-rates-for-india-polls-ads-than-others

[22] A. Kaul and D. Kumar, "Tek Fog: An App With BJP Footprints for Cyber Troops to Automate Hate, Manipulate Trends," The Wire, Jan. 6, 2022. https://thewire.in/tekfog/en/1.html.

[23] ACM FAccT's own Strategic Plan lays out a number of admirable principles that could have laid the groundwork for just such a vision. "Our conference should feature tremendous diversity in the computational systems studied, the sociotechnical problems considered, and the approaches to understanding, and mitigation employed." (see https://facctconference.org/static/docs/strategic_plan.pdf). This statement attests to the importance of multiple voices and approaches to problem solving. It goes on to express a commitment to the situated social contexts in which systems operate: "Work that does not have deep engagement with the social component of problems … [is] considered outside the bounds of the conference." It also acknowledges the seriousness of the stakes with which the conference contends, in arguing that FAccT "Establish

in Global South/Global Majority settings.[24] While a few advocates and activists—particularly from large or well-funded organizations—have been present as attendees or session organizers, they are rarely present as distinguished guests or keynote speakers. They also are seldom from Global Majority contexts, smaller and low-resource organizations, or without institutional affiliations. FAccT's Critiquing and Rethinking Accountability, Fairness and Transparency (CRAFT) track was founded in part to open the conference to a wider array of contributions, but we find that the CRAFT track off-and-to-the-side of the main conference does not reflect its centrality to FAccT's stated aims. The primarily academic valence of the conference leaves the opportunity to approach AI ethics from a wider lens unrealized. We call on the FAccT organizing committee to conduct formal outreach to movement activists and advocates, such as through paid focus groups, to examine what the conference might be able to offer—in particular how the conference could be be aligned with activists' and advocates' theories of social change.

4.2.1 2022 Social Media Influencers Conference as precedent. We look to other conferences which have re-imagined what a conference in our field might be, even as it takes place within an academic institution. On April 7-8 2022, the Social Media Influencers and the New Political Economy in South Asia and Africa <influencers.conference.si.umich.edu> organized by Professors Joyojeet Pal and Omolade Adunbi convened activists, artists, politicians, dissidents, industry practitioners, comedians, journalists, actors, NGO leaders, entrepreneurs, and media figures from India, Sri Lanka, Nigeria, Kenya, and Pakistan as guests and invited contributors. Academic participation was woven into the practitioner-focused program. All speakers drew attention to the uses and consequences of tech platforms on the ground. The conversations possible across these different vantage points fostered ties across national contexts as well as exchange of tactics and experiences. Rather than be thrust into a primarily academic event in a vestigial capacity, advocate and activist experiences were centered within the space—acknowledging the primacy of situated knowing in questions of justice and power.

We argue that the Social Media Influencers Conference (SMIC) is a model example of what it would mean for an academic conference to share space with advocates and activists. The urgent questions that emerged in this forum are visible to interested readers under the #SMIC2022 hashtag; they include powerful testimony from Gulalai Ismail (Gulalai_Ismail), a Pashtun human rights activist, who on February 6, 2019 was disappeared by the Pakistani government from a protest of extrajudicial killing and the murder of activist Arman Loni earlier that month. She attributes mass movement organizing on social media via #wheresgulalai to her release and safe return. In the Pakistani social media context, she notes how algorithmic systems' ability to boost her story into machine learning-driven top line news under the "Trending" tab led to her safe release; at the same time, this engine of mass dissemination based on user engagement caused a video of girls dancing to result in their being targeted for murder in the 2012 Kohistan case. Ismail's remarks at the conference called on technology companies to be more culturally-sensitive and nuanced—not relativist, but instead to localize and democratize their services to those who understand cultural context and consequences. The SMIC conference also featured investigative journalist Josy Joseph (josyjosephkj) who told the audience that academics play an important role in confronting powerful interests in country contexts where investigative journalists are prosecuted. In particular, he highlighted the role that social media companies play in amplifying harm through hate speech, saying "Many social media platforms are human rights offenders, killers at large scale." We believe that FAccT could present another key venue for experiences and perspectives like Ismail's and Joseph's to be put in conversation with other FAccT audiences, including activists working on other social issues, tech justice activists, and individuals with institutional power relating to AI policy and design decisions about safety and harm.

4.2.2 2020 Resistance AI Workshop as precedent. Another venue with an inclusive approach was the 2020 Resistance AI Workshop that occurred on December 11, 2020, hosted by the Radical AI Network at the Neural Information Processing Systems (NeurIPS) 2020 Conference. Resistance AI articulated the question of power in its mission:

> It has become increasingly clear in recent years that AI research, far from producing neutral tools, has been concentrating power in the hands of governments and companies and away from marginalized communities. Unfortunately, the NeurIPS Conference—one of the largest and most esteemed machine learning conferences in the world—has until now lacked a venue explicitly dedicated to understanding and addressing this concerning reality. As Black feminist scholar Angela Davis famously said, "Radical simply means grasping things at the root." Resistance AI exposes the root of the current reality: technology rearranges power. We believe that when we are engaged in Resistance AI, we can both resist AI that centralizes power in the hands of the few and we can dream up and build human/AI systems that put power in the hands of the people. This workshop will be a space for AI researchers and marginalized communities to discuss and reflect on AI-fueled inequity and co-create our dreams and tactics of how to work toward Resistance AI.

This workshop accepted contributions in a variety of formats and languages, from an infographic,[25] to a poem,[26] to a comic,[27] to

and maintain a culture of modesty out of respect for the gravity of the problems that are the focus of our community."

[24] See R. Singh and R. Lara Guzmán, "Parables of AI infrom the Global South," Society for Social Studies of Science, Sept. 17, 2021. https://www.4sonline.org/parables-of-ai-in-from-the-global-south/.

[25] C. D'Ignazio, L. Klein, M. Diaz, "The Data Feminism Infographic," https://drive.google.com/file/d/1EHnQ34wB1PMo_lZX2sBbLrhAj-XkKddL/view.

[26] H. Jethwani, "Panopticon (Resistance Poem)" https://drive.google.com/file/d/1A6j_ZfLM6P91V-_1SQdw5UxcgqMzVedc/view.

[27] F. A. Khan, J. Stoyanovich, "Mirror, Mirror," https://drive.google.com/file/d/1IMD51ih4-9C-CoATCwIFT5CKWXDz42e-/view.

speculative fiction,[28] to sound art,[29] and an illustrated parody[30] of Google-sponsored research.

4.2.3 Going beyond inclusion. Inclusion of people formerly excluded from predominantly white, elite, or Global North institutions is not a sufficient remediation for centuries of discrimination, which are reproduced by academic and commercial technology innovation. The "stakeholder turn" in HCI and related disciplines once demanded acknowledgement that many people have a stake in technological design [21]. Despite the aims of this work to democratize and widen the scope of perspective, stakeholder language presumes equality of power among participants to a debate or consultation [28, 35]. Including a diverse group of people in the process does not guarantee improved conditions for those most marginalized; simply having a "seat at the table" is insufficient if the table is dominated by participants with power, influence, and presumptions of epistemic superiority.

4.3 Embracing dissensus: An agonistic approach

We find conflict of interest risk is tied to the current institutional design of the FAccT conference. Here we re-imagine FAccT in an agonistic vein. Agonism refers to the positive or generative aspects of conflict, drawing on Belgian political theorist Chantal Mouffe, whose "agonistic pluralism" is founded on the idea that a certain amount of conflict is intrinsic to political matters and differently situated people and interests [49]. Mouffe argues that foreclosing this conflict through false consensus is always power-laden and a frontier of inclusion and exclusion [37]. Instead, she embraces an approach that would mobilize the full plurality of conflicting views toward a more democratic whole. The promise of this approach is to empower and platform formerly excluded people and to directly address harms of marginalization [20]. In the context of FAccT, we have argued that the conference proceedings produces a similarly power-laden image of consensus that ultimately fails to shift power relations in a meaningful way. Re-imagining the FAccT conference in an agonistic vein means opening up space for the full constellation of divergent interests and opinions that AI ethics implicates, for instance to:

- Change the format of the conference program to feature explicit debates about how to advance the aims of fairness, accountability, and transparency.
- Fold CRAFT workshops and interventions into the body of the conference itself, rather than bracketing critical and participatory contributions as something aside from its primary aims.
- Invite keynote speakers who can intervene directly on the audience as a group of people ourselves to be organized.
- Expand the scope of the Town Hall segment, traditionally at the end of the conference, to include themes such as techniques for funding disclosure, managing conflicts of interest, holding space for sharing personal experiences, and reconfiguring the conference program to better meet its stated aims.
- Follow the lead of the most marginalized people in our societies rather than positioning ourselves to make positive social change on behalf of those most marginalized (Cf., [23, 28]).

We imagine these changes will foster a more active and lively conference that moves beyond a passive presentation of findings to a forum in which participants come to recognize and acknowledge the political import of the field [42]. This shift would also move the conference further toward FAccT's governance documents, which state a commitment to be "an independent conference not beholden to anyone."[31]

5 CONCLUSION

FAccT has emerged as the preeminent professional body for facilitating dialogue on AI ethics and governance, granting it a position of significant power and prestige. Yet rather than using this position to hold space for people on the receiving end of algorithmic harm, FAccT instead serves power, particularly the technology firms that underwrite the conference and participate in much of its work. Firms take up significant space in this forum and in its proceedings. In the midst of a crisis of trust and integrity in the field of ML and AI, FAccT has met the needs of technology firms while failing to meet the needs of real people facing significant, in some cases life-threatening, struggles exacerbated or caused by technology platforms [2]. Here we quote Birhane et. al [11] who measure the scale of this neglect across four years of proceedings in FAccT and the closely related AI, Ethics, and Society conference:

> A review of the AI ethics papers from two of the most prominent conferences showed there is a great tendency for abstract discussion of such topics devoid of structural and social factors and specific and potential harms. Given that the most marginalized in society are the most impacted when algorithmic systems fail, we contend that all AI ethics work, from research, to policy, to governance, should pay attention to structural factors and actual existing harms.

Academic AI ethics conferences such as ACM FAccT and other venues of authoritative knowledge production in computing must take greater responsibility for the ideas they spread and the credibility they lend—particularly when those ideas and works reach beyond disciplinary boundaries to broader circles of influence. The urgency of social and ecological crises facing the world demand a reconfiguration of institutions of knowledge production whose status quo operations serve to reproduce social, economic, and political injustice.

In this paper, we have argued that FAccT must adopt rigorous financial disclosure policies for its authors and organizers. We have also argued that FAccT must do significantly more to become a

[28] N. T. Djanegara, "Illegible,"https://ninadewi.github.io/Resistance%20AI_Illegible.pdf.
[29] D. Sinha, "Dhakuria Bridge" (description at https://drive.google.com/file/d/1bebkl453Zy-XGnLg1emVlYSkeu539MFB/view) and audio at https://soundcloud.com/debsinha/sets/dhakuria-bridge/s-KNJzpJ0GgVy).
[30] Critical Platform Studies Group, "The A-Z of UAVs," https://archive.org/details/the-a-z-of-uavs.
[31] "ACM FAccT Strategic Plan," ACM FAccT Conference https://facctconference.org/static/docs/strategic_plan.pdf.

welcoming and inviting space for a broader range of scholars and activists, such as from small, low-resource, grassroots, community-oriented, Global South/Global Majority, or BIPOC-led organizations as well as individuals without social capital or institutional affiliations. By doing so, the FAccT community is far more likely to feature perspectives on algorithmic harm that challenge the products and practices of firms, who are already well-represented at the conference.

FAccT must also respond to contemporary and historical struggles against oppression by engaging in critical reflection and by re-centering its organizing principles and work [15, 17, 34, 70] toward political action [27, 48, 65], resistance [30, 68], and justice [14]. Venues such as FAccT would better serve the aims of AI ethics by embracing an "agonistic" approach to discourse, both within its research subjects and in its organizing principles. This approach is more likely to produce institutions and work that focus on fundamental questions of fairness, transparency, and accountability. This may require decentering technology as a main nominal focus in favor of substantive issues that would build stronger bridges to advocates and activists fighting for social, economic, and environmental justice [53].

FUNDING DISCLOSURE

The first author received no further funding aside from the salary provided by her institution drawn from award #1801307 on "Contextual Integrity: From Theory to Practice" from the National Science Foundation. She participated in a Microsoft Research internship in the past two years. The second author received no further funding aside from his salary, which is supported in part by a non-departmental public body of the Government of the United Kingdom called UK Research & Innovation via its Economic and Social Research Council Grant ES/T007354/1, Engineering and Physical Sciences Research Council (EPSRC) Grant EP/T001569/1, EPSRC Grant EP/W006022/1, and from the public funds that make the Turing's Public Policy Programme possible. The third author received no further funding aside from their salary, 1/6 of which is covered by a November 2021 grant from UK Research & Innovation called "Transforming Collections: Reimagining Art, Nation and Heritage;" the rest is drawn from a general fund at their institution into which they have limited visibility but which is composed primarily of student fees.

The following disclosure accords with the detailed disclosure guidelines in FAccT's new Conflict of Interest policy.[32] None of the authors' employers or funders played a role in designing or conducting this work. None of the authors' employers or funders played a role in accessing or collecting data. None of the authors' employers or funders played a role in data analysis or interpretation. None of the authors' employers or funders played a role in preparing, reviewing, or approving the manuscript. None of the authors' employers or funders played a role in deciding to submit the manuscript. The authors have no financial ties to any entity affected by publication of the paper. They acknowledge a social relationship to the founder of Parity, which is a competitor of Pymetrics, a firm criticized at length in our work.

[32] "Disclosure of Funding Sources at ACM FAccT 2022" ACM FAccT Conference https://facctconference.org/2022/funding_sources_disclosure.html

ACKNOWLEDGMENTS

This piece is indebted to the co-organizers and signatories of the 2021 letter to FAccT's Executive Committee (EC) calling on the EC to change FAccT's direction, public perception, and priorities [36]; as well as participants in conversations around that letter, including especially J Khadijah Abdurahman and Seda Gürses. We would also like to specifically acknowledge Anna Lauren Hoffmann, whose scholarship on refusal and the limits of inclusion has greatly furthered our thinking. Thank you to Anmol Panda for his numerous contributions to the examples we draw from his area of expertise. Thank you to Joyojeet Pal and Omolade Adunbi for discussions on their work to reimagine academic gatherings through their 2022 Social Media Influencers Conference. We thank Emanuel Moss for sharing his manuscript on the history and evolution of the FAccT conference on background, Thomas Gilbert for the perspective and sources he shared on agonism and antagonism, and Seth Pollack for reviewing our draft. Thank you to M Govindu for making the point about the relationship between exclusion and paternalism. Thank you to past panelists, anonymous contributors, and organizers in our Patron or Poison event and campaign on industry funding of HCI. We would also like to thank our anonymous reviewers for the profoundly insightful and developmental feedback on our first draft; their recommended readings and incisive questions immeasurably improved the quality of this work. Finally, thank you to those who have donated to the Critical Platform Studies Group. You can support our work at https://www.patreon.com/critplat. This document is an expression of one moment in an ongoing process of learning; all errors and omissions are our own.

REFERENCES

[1] Mohamed Abdalla and Moustafa Abdalla. 2021. The Grey Hoodie Project: Big Tobacco, Big Tech, and the threat on academic integrity. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 287–297.

[2] J. Khadijah Abdurahman. 2020. On the moral collapse of AI ethics. https://upfromthecracks.medium.com/on-the-moral-collapse-of-ai-ethics-791cbc7df872.

[3] Sara Ahmed. 2012. On Being Included. In *On Being Included*. Duke University Press. https://doi.org/10.2307/j.ctv1131d2g

[4] Syeda Zainab Akbar, Ankur Sharma, Himani Negi, Anmol Panda, and Joyojeet Pal. 2020. Anatomy of a rumour: Social media and the suicide of Sushant Singh Rajput. *arXiv preprint arXiv:2009.11744* (2020).

[5] Bodil Als-Nielsen, Wendong Chen, Christian Gluud, and Lise L Kjaergard. 2003. Association of funding and conclusions in randomized drug trials: A reflection of treatment effect or adverse events? *JAMA* 290, 7 (2003), 921–928. https://doi.org/10.1001/jama.290.7.921

[6] Amarnath Amarasingam, Sanober Umar, and Shweta Desai. 2022. "Fight, die, and if required kill": Hindu nationalism, misinformation, and Islamophobia in India. *Religions* 13, 5 (2022), 380. https://doi.org/10.3390/rel13050380

[7] Justin E Bekelman, Yan Li, and Cary P Gross. 2003. Scope and impact of financial conflicts of interest in biomedical research: A systematic review. *JAMA* 289, 4 (2003), 454–465. https://doi.org/10.1001/jama.289.4.454

[8] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 610–623. https://doi.org/10.1145/3442188.3445922

[9] Elettra Bietti. 2020. From ethics washing to ethics bashing: A view on tech ethics from within moral philosophy. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 210–219.

[10] Reuben Binns. 2018. Algorithmic accountability and public reason. *Philosophy & Technology* 31, 4 (2018), 543–556. https://doi.org/10.1007/s13347-017-0263-5

[11] Abeba Birhane, Elayne Ruane, Thomas Laurent, Matthew Brown, Johnathan Flowers, Anthony Ventresque, and Chistopher Dancy. 2022. The forgotten margins of AI ethics. In *Proceedings of the 2022 Conference on Fairness, Accountability, and Transparency*.

[12] Parmit K Chilana, Mary P Czerwinski, Tovi Grossman, Chris Harrison, Ranjitha Kumar, Tapan S Parikh, and Shumin Zhai. 2015. Technology transfer of HCI

research innovations: Challenges and opportunities. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*. 823–828. https://doi.org/10.1145/2702613.2724724
[13] Sasha Costanza-Chock. 2020. *Design Justice: Community-Led Practices to Build the Worlds We Need*. The MIT Press. https://doi.org/10.7551/mitpress/12255.001.0001
[14] Lina Dencik. 2021. Towards data justice unionism? A labour perspective on AI governance. *AI for Everyone* (2021), 267–284. https://doi.org/10.16997/book55.o
[15] Lina Dencik, Arne Hintz, Joanna Redden, and Emiliano Treré. 2019. Exploring data justice: Conceptions, applications and directions. , 873–881 pages. https://doi.org/10.1080/1369118X.2019.1606268
[16] Roderick A Ferguson. 2012. *The Reorder of Things: The University and its Pedagogies of Minority Difference*. University of Minnesota Press.
[17] Rodrigo Ferreira and Moshe Y Vardi. 2021. Deep tech ethics: An approach to teaching social justice in computer science. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*. 1041–1047. https://doi.org/10.1145/3408877.3432449
[18] Marilyn J Field, Bernard Lo, et al. 2009. Conflict of interest in medical research, education, and practice. (2009).
[19] Thomas E Finucane and Chad E Boult. 2004. Association of funding and findings of pharmaceutical research at a meeting of a medical professional society. *The American journal of medicine* 117, 11 (2004), 842–845. https://doi.org/10.1016/j.amjmed.2004.05.029
[20] Thomas Fossen. 2008. Agonistic critiques of liberalism: Perfection and emancipation. *Contemporary Political Theory* 7, 4 (2008), 376–394.
[21] Batya Friedman and David Hendry. 2019. *Value Sensitive Design: Shaping Technology with Moral Imagination*. MIT Press.
[22] James F Fries and Eswar Krishnan. 2004. Equipoise, design bias, and randomized controlled trials: The elusive ethics of new drug development. *Arthritis Res Ther* 6, 3 (2004), 1–6. https://doi.org/10.1186/ar1170
[23] Maya Indira Ganesh and Emanuel Moss. 2022. Resistance and refusal to algorithmic harms: Varieties of 'knowledge projects'. *Media International Australia* 183, 1 (2022), 90–106. https://doi.org/10.1177/1329878X221076288
[24] Ben Gansky and Sean McDonald. 2022. CounterFAccTual: How FAccT undermines its organizing principles. In *Proceedings of the 2022 Conference on Fairness, Accountability, and Transparency*.
[25] Kiran Garimella and Dean Eckles. 2020. Images and misinformation in political groups: Evidence from WhatsApp in India. *arXiv preprint arXiv:2005.09784* (2020).
[26] Ben Green. 2021. The contestation of tech ethics: A sociotechnical approach to technology ethics in practice. *Journal of Social Computing* 2, 3 (2021), 209–225. https://doi.org/10.23919/JSC.2021.0018
[27] Ben Green. 2021. Data science as political action: Grounding data science in a politics of justice. *Journal of Social Computing* 2, 3 (2021), 249–265. https://doi.org/10.23919/JSC.2021.0029
[28] Daniel Greene, Anna Lauren Hoffmann, and Luke Stark. 2019. Better, nicer, clearer, fairer: A critical assessment of the movement for ethical artificial intelligence and machine learning. In *Proceedings of the 52nd Hawaii international conference on system sciences*.
[29] Critical Platform Studies Group, Lilly Irani, Niloufar Salehi, Joyojeet Pal, Andrés Monroy-Hernández, Elizabeth Churchill, and Sneha Narayan. 2019. Patron or Poison? Industry funding of HCI research. In *Conference Companion Publication of the 2019 on Computer Supported Cooperative Work and Social Computing*. 111–115.
[30] Thilo Hagendorff. 2021. Blind spots in AI ethics. *AI and Ethics* (2021), 1–17. https://doi.org/10.1007/s43681-021-00122-8
[31] Alex Hanna. 2022. On racialized tech organizations and complaint: A goodbye to Google. https://alex-hanna.medium.com/on-racialized-tech-organizations-and-complaint-a-goodbye-to-google-43fd8045991d.
[32] Alex Hanna and Emily Denton. 2020. CtrlZ.AI Zine Fair: Critical perspectives. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) *(FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 686. https://doi.org/10.1145/3351095.3375692
[33] Stefano Harney and Fred Moten. 2013. *The Undercommons: Fugitive Planning and Black Study*. Minor Compositions.
[34] Anna Lauren Hoffmann. 2019. Where fairness fails: Data, algorithms, and the limits of antidiscrimination discourse. *Information, Communication & Society* 22, 7 (2019), 900–915. https://doi.org/10.1080/1369118X.2019.1573912
[35] Anna Lauren Hoffmann. 2021. Terms of inclusion: Data, discourse, violence. *New Media & Society* 23, 12 (2021), 3539–3556. https://doi.org/10.1177/1461444820958725
[36] Anna Lauren Hoffmann, Lilly Irani, Michael Katell, P M Krafft, Vinhcent Le, Dan Greene, Meredith Whitaker, Niloufar Salehi, Meg Young, Linnet Taylor, Michael Veale, Catherine D'Ignazio, Jevan Hutson, Jennifer Lee, Reuben Binns, Luke Stark, Maya Indira Ganesh, Alex Hanna, and Ben Green. 2021. FAccT conference - direction, perception, and priorities. Private communication to FAccT Executive Committee, 9 Dec 2021.
[37] Matthew Jones. 2014. Chantal Mouffe's agonistic project: Passions and participation. *Parallax* 20, 2 (2014), 14–30.
[38] Robert E Kelly, Lisa J Cohen, Randye J Semple, Philip Bialer, Adam Lau, Alison Bodenheimer, Elana Neustadter, Arkady Barenboim, and Igor I Galynker. 2006. Relationship between drug company funding and outcomes of clinical psychiatric research. *Psychological medicine* 36, 11 (2006), 1647–1656. https://doi.org/10.1017/S0033291706008567
[39] Issa Kohler-Hausmann. 2018. Eddie Murphy and the dangers of counterfactual causal thinking about detecting racial discrimination. *Northwestern University Law Review* 113 (2018), 1163.
[40] la paperson. 2017. *A Third University is Possible*. University of Minnesota Press. https://doi.org/10.1017/S0007680512001031
[41] Benjamin Laufer, Sameer Jain, A. Feder Cooper, Jon Kleinberg, and Hoda Heidari. 2022. Four years of FAccT: A reflexive, mixed-methods analysis of research contributions, shortcomings, and future prospects. In *Proceedings of the 2022 Conference on Fairness, Accountability, and Transparency*.
[42] Jennifer Lee, Meg Young, PM Krafft, and Michael A Katell. 2020. Power and technology: who gets to make the decisions? *Interactions* 28, 1 (2020), 38–46. https://doi.org/10.1145/3442420
[43] Joel Lexchin, Lisa A Bero, Benjamin Djulbegovic, and Otavio Clark. 2003. Pharmaceutical industry sponsorship and research outcome and quality: systematic review. *BMJ* 326, 7400 (2003), 1167–1170. https://doi.org/10.1136/bmj.326.7400.1167
[44] Howard Liss. 2006. Publication bias in the pulmonary/allergy literature: Effect of pharmaceutical company sponsorship. *Focus* (2006).
[45] Donald McMillan and Barry Brown. 2019. Against ethical AI. In *Proceedings of the Halfway to the Future Symposium*. 1–3.
[46] Jodi Melamed. 2016. Being together subversively, outside in the university of hegemonic affirmation and repressive violence, as things heat up (again). *American Quarterly* 68, 4 (2016), 981–991.
[47] Philip Mirowski. 2011. Science-Mart. In *Science-Mart*. Harvard University Press. https://doi.org/10.4159/harvard.9780674061132
[48] Jared Moore. 2020. Towards a more representative politics in the ethics of computer science. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 414–424. https://doi.org/10.1145/3351095.3372854
[49] Chantal Mouffe. 2013. *Agonistics: Thinking the World Politically*. Verso Books.
[50] Rodrigo Ochigame. 2019. The invention of "ethical AI": How Big Tech manipulates academia to avoid regulation. *The Intercept* (2019). https://theintercept.com/2019/12/20/mit-ethical-ai-artificial-intelligence/.
[51] Naomi Oreskes and Erik M Conway. 2011. *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco smoke to Global Warming*. Bloomsbury Publishing USA.
[52] Anmol Panda, Sunandan Chakraborty, Noopur Raval, Han Zhang, Mugdha Mohapatra, Syeda Zainab Akbar, and Joyojeet Pal. 2020. Affording Extremes: Incivility, Social Media and Democracy in the Indian Context. In *Proceedings of the 2020 International Conference on Information and Communication Technologies and Development* (Guayaquil, Ecuador) *(ICTD2020)*. Association for Computing Machinery, New York, NY, USA, Article 9, 12 pages. https://doi.org/10.1145/3392561.3394637
[53] Seeta Peña Gangadharan and Jędrzej Niklas. 2019. Decentering technology in discourse on discrimination. *Information, Communication & Society* 22, 7 (2019), 882–899. https://doi.org/10.1080/1369118X.2019.1593484
[54] Inioluwa Deborah Raji, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) *(FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 33–44. https://doi.org/10.1145/3351095.3372873
[55] Paul M Ridker and Jose Torres. 2006. Reported outcomes in major cardiovascular clinical trials funded by for-profit and not-for-profit organizations: 2000-2005. *JAMA* 295, 19 (2006), 2270–2274. https://doi.org/10.1001/jama.295.19.2270
[56] Paula A Rochon, Jerry H Gurwitz, Robert W Simms, Paul R Fortin, David T Felson, Kenneth L Minaker, and Thomas C Chalmers. 1994. A study of manufacturer-supported trials of nonsteroidal anti-inflammatory drugs in the treatment of arthritis. *Archives of Internal Medicine* 154, 2 (1994), 157–163. https://doi.org/10.1001/archinte.1994.00420020059007
[57] Steven Shapin. 2009. The Scientific Life. In *The Scientific Life*. University of Chicago Press.
[58] Hong Shen, Alicia DeVos, Motahhare Eslami, and Kenneth Holstein. 2021. Everyday algorithm auditing: Understanding the power of everyday users in surfacing harmful algorithmic behaviors. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–29. https://doi.org/10.1145/3479577
[59] Sergio Sismondo. 2008. Pharmaceutical company funding and its consequences: A qualitative systematic review. *Contemporary Clinical Trials* 29, 2 (2008), 109–113. https://doi.org/10.1016/j.cct.2007.08.001
[60] Mona Sloane. 2019. Inequality is the name of the game: thoughts on the emerging field of technology, ethics and social justice. In *Weizenbaum Conference*. DEU, 9.
[61] Constantin Vica, Cristina Voinea, and Radu Uszkai. 2021. The emperor is naked: Moral diplomacies and the ethics of AI. *Információs Társadalom* 21, 2 (2021).
[62] Lee Vinsel. 2012. Review of Science-Mart: Privatizing American science. *Business History Review* 86, 3 (2012), 650–653.
[63] Sandra Wachter, Brent Mittelstadt, and Luciano Floridi. 2017. Why a right to explanation of automated decision-making does not exist in the general data

protection regulation. *International Data Privacy Law* 7, 2 (2017), 76–99. https://doi.org/10.1093/idpl/ipx005

[64] Ben Wagner. 2018. Ethics As An Escape From Regulation. From "Ethics-Washing" To Ethics-Shopping? In *BEING PROFILED:COGITAS ERGO SUM: COGITAS ERGO SUM: 10 Years of Profiling the European Citizen*, Emre Bayamlioglu, Irina Baraliuc, Liisa Albertha Wilhelmina Janssens, and Mireille Hildebrandt (Eds.). Amsterdam University Press, 84–89. https://doi.org/doi:10.1515/9789048550180-016

[65] Meredith Whittaker. 2021. The steep cost of capture. *Interactions* 28, 6 (2021), 50–55. https://doi.org/10.1145/3488666

[66] Oscar Williams. 2019. How Big Tech funds the debate on AI ethics. *The New Statesman* (2019).

[67] Christo Wilson, Avijit Ghosh, Shan Jiang, Alan Mislove, Lewis Baker, Janelle Szary, Kelly Trindel, and Frida Polli. 2021. Building and auditing fair algorithms: A case study in candidate screening. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. Association for Computing Machinery, 666–677. https://doi.org/10.1145/3442188.3445928

[68] Resistance AI Workshop. 2020. Why we are hosting the Resistance AI Workshop. *NeurIPS* (2020). https://sites.google.com/view/resistance-ai-neurips-20/home.

[69] Meg Young, Luke Rodriguez, Emily Keller, Feiyang Sun, Boyang Sa, Jan Whittington, and Bill Howe. 2019. Beyond open vs. closed: Balancing individual privacy and public accountability in data sharing. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 191–200. https://doi.org/10.1145/3287560.3287577

[70] Andrea Zeffiro. 2021. From data ethics to data justice in/as pedagogy (Dispatch). *Studies in Social Justice* 15, 3 (2021), 450–457. https://doi.org/10.26522/ssj.v15i3.2546